

# Efficient Dense Gaussian Elimination over the Finite Field with Two Elements

MARTIN R. ALBRECHT

INRIA, Paris-Rocquencourt Center, SALSA Project UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France; CNRS, UMR 7606, LIP6, F-75005, Paris, France

GREGORY V. BARD

University of Wisconsin-Stout, Menomonie, Wisconsin 54751, USA

and

CLÉMENT PERNET

INRIA-MOAIS LIG, Grenoble Univ. ENSIMAG, Antenne de Montbonnot 51, avenue Jean Kuntzmann, F-38330 MONTBONNOT SAINT MARTIN, France

---

In this work we describe an efficient implementation of a hierarchy of algorithms for Gaussian elimination upon dense matrices over the field with two elements ( $\mathbb{F}_2$ ). We discuss both well-known and new algorithms as well as our implementations in the M4RI library, which has been adopted into SAGE. The focus of our discussion is a block iterative algorithm for PLE decomposition which is inspired by the M4RI algorithm. The implementation presented in this work provides considerable performance gains in practice when compared to the previously fastest implementation. We provide performance figures on x86\_64 CPUs to demonstrate the alacrity of our approach.

Categories and Subject Descriptors: G.4 [MATHEMATICAL SOFTWARE]:

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: GF(2), RREF, Linear Algebra, Gaussian Elimination, Matrix Decomposition, Method of Four Russians

---

## 1. INTRODUCTION

We describe an efficient implementation of a hierarchy of algorithms for Gaussian elimination of dense matrices over the field with two elements ( $\mathbb{F}_2$ ). In particular, we describe algorithms for PLE decomposition [Jeannerod et al. 2011] which decompose a matrix  $A$  into  $P \cdot L \cdot E$  where  $P$  is a permutation matrix,  $L$  is a unit lower triangular matrix and  $E$  a matrix in row echelon form. Since matrix decomposition is an essential building block for solving dense systems of linear and non-linear equations [Lazard 1983] much research has been devoted to improving the asymptotic complexity of such algorithms. A line of research which produced decompositions such as PLUQ [Golub and Van Loan 1996], LQUP [Ibarra et al. 1982], LSP [Ibarra et al. 1982] and PLE [Jeannerod et al. 2011], among others. Each of those decompositions can be reduced to matrix-matrix multiplication and hence can be achieved in  $\mathcal{O}(n^\omega)$  time where  $\omega$  is the complexity exponent of linear algebra<sup>1</sup>. However, in [Jeannerod et al. 2011] it was shown that many of these

---

<sup>1</sup>For practical purposes we set  $\omega = 2.807$ . Lower exponents have been obtained in theory [Coppersmith and Winograd 1990].

decompositions are essentially equivalent (one can be produced from the other in negligible cost) and that for almost all applications the PLE decomposition is at least as efficient in terms of time and memory as any of the other. Hence, in this work, we focus on the PLE decomposition but consider the special case of  $\mathbb{F}_2$ . In particular, we propose a new algorithm for block-iterative PLE decomposition and discuss its implementation. We also describe our implementation of previously known algorithms in the M4RI library [Albrecht and Bard 2011].

This article is organised as follows. We will first highlight some computer architecture issues and discuss in place methods, as well as define notation in the remainder of this section. We will start by giving the definitions of reduced row echelon forms (RREF) and PLE decomposition in Section 2. We will then discuss the three algorithms and their implementation issues for performing PLE decomposition in Section 3, and other implementation issues in Section 4. We conclude by giving empirical evidence of the viability of our approach in Section 5. In Appendix A we will describe the original M4RI algorithm (for reducing  $\mathbb{F}_2$  matrices into either row echelon form or RREF), which initiated this line of research.

### 1.1 Computer Architecture Issues

The M4RI library implements dense linear algebra over  $\mathbb{F}_2$ . Our implementation focuses on 64-bit x86 architectures (x86\_64), specifically the Intel Core 2 and the AMD Opteron. Thus, we assume in this work that native CPU words have 64 bits. However, it should be noted that our code also runs on 32-bit CPUs and on non-x86 CPUs such as the PowerPC.

Element-wise operations over  $\mathbb{F}_2$ , being mathematically trivial, are relatively cheap compared to memory access. In fact, in this work we demonstrate that the two fastest implementations for dense matrix decomposition over  $\mathbb{F}_2$  (the one presented in this work and the one found in MAGMA [Bosma et al. 1997] due to Allan Steel) perform worse for moderately sparse matrices despite the fact that fewer field operations are performed. Thus our work provides further evidence that more refined models for estimating the expected efficiency of linear algebra algorithms on modern CPUs and their memory architectures are needed.

An example of the CPU-specific issues would be the “SSE2 instruction set,” which we will have cause to mention in Section 4. This instruction set has a command to XOR (exclusive-OR) two 128-bit vectors, with a single instruction. Thus if we are adding a pair of rows of  $128n$  bits, we can do that with  $n$  instructions. This is a tremendous improvement on the  $128n$  floating-point operations required to do that with real-numbered problems. This also means that any algorithm which touches individual bits instead of blocks of 128 bits will run about two orders of magnitude slower than an algorithm which avoids this.

### 1.2 Notations and space efficient representations

By  $A[i, j]$  we denote the entry of the matrix  $A$  at row  $i$  and column  $j$ . By  $\ell_i$  we denote the  $i$ th element of the vector  $\ell$ . We represent  $m \times m$  permutation matrices as integer vectors of length  $m$ , i.e., in LAPACK-style. That is to say, the permutation

matrix

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

is stored as  $P = [0, 2, 2, 4, 4]$ , where for each index  $i$  the entry  $P_i$  encodes which swap ought to be performed on the input. This representation allows to apply permutations in-place.

All indices start counting at zero, i.e.  $A[2, 2]$  refers to the intersection of the third row and third column. All algorithms, lemmas, theorems, propositions in this work are presented for the special case of  $\mathbb{F}_2$ . Hence, any statement that some algorithm produces a certain output is always understood to apply to the special case of  $\mathbb{F}_2$ , unless stated otherwise.

Matrix triangular decompositions allow that the lower triangular and the upper triangular matrices can be stored one above the other in the same amount of memory as for the input matrix: their non trivial coefficients occupy disjoint areas, as the diagonal of one of them has to be unit and can therefore be omitted. For two  $m \times n$  matrices  $L$  and  $U$  such that  $L = [\ell_{i,j}]$  is unit lower triangular and  $U = [u_{i,j}]$  is upper triangular, we shall use the notation  $L \setminus U$ , following [Jeannerod et al. 2011] to express the fact that  $L$  and  $U$  are stored one next to the other within the same  $m \times n$  matrix. Thus,  $A = L \setminus U$  is the  $m \times n$  matrix  $A = [a_{i,j}]$  such that  $a_{i,j} = \ell_{i,j}$  if  $i > j$ , and  $a_{i,j} = u_{i,j}$  otherwise. We call it a space sharing representation of  $L$  and  $U$ . Furthermore,  $A = A_0 \mid A_1$  means that  $A$  is obtained by augmenting  $A_0$  by  $A_1$ .

The methods of this paper not only produce space-sharing triangular decompositions, but also overwrite their input matrix with the output triangular decomposition. However, they are not in-place, contrarily to the general algorithms in [Jeannerod et al. 2011], as they use tabulated multiplications that require non constant extra memory allocations.

Last but not least, complexity is always reported for square matrices, but the complexities for over-defined and under-defined matrices can be derived with moderate effort.

## 2. MATRIX DECOMPOSITIONS

In this work we are interested in matrix decompositions that do not destroy the column profile, since those play a crucial role in linear algebra for Gröbner basis computations [Faugère 1999]<sup>2</sup>.

**DEFINITION 1 COLUMN RANK PROFILE.** *The column rank profile of a matrix  $A$  is the lexicographically smallest sequence of column indices  $i_0 \leq \dots < i_r$ , with  $r$  the rank of  $A$  such that the corresponding columns are linearly independent.*

**DEFINITION 2 LEADING ENTRY.** *The leading entry in a non-zero row of a matrix is the leftmost non-zero entry in that row.*

<sup>2</sup>These are also useful in simplified variants of the F4 algorithm for polynomial system solving such as the popular [Courtois et al. 2000].

Perhaps, the most well-known of such decompositions is the row echelon form.

**DEFINITION 3 ROW ECHELON FORM.** *An  $m \times n$  matrix  $A$  is in row echelon form if any all-zero rows are grouped together at the bottom of the matrix, and if the leading entry of each non-zero row is located to the right of the leading entry of the previous row.*

While row echelon forms are not unique, *reduced* row echelon forms, whose definition we reproduce below, are.

**DEFINITION 4 REDUCED ROW ECHELON FORM.** *An  $m \times n$  matrix  $A$  is in reduced row echelon form if it is in row echelon form, and furthermore, each leading entry of a non-zero row is the only non-zero element in its column.*

In order to compute (reduced) row echelon forms, we compute the *PLE* decomposition.

**DEFINITION 5 PLE DECOMPOSITION.** *Let  $A$  be an  $m \times n$  matrix over a field  $\mathbb{F}$ . A PLE decomposition of  $A$  is a triple of matrices  $P, L$  and  $E$  such that  $P$  is an  $m \times m$  permutation matrix,  $L$  is an  $m \times m$  unit lower triangular matrix, and  $E$  is an  $m \times n$  matrix in row-echelon form, and  $A = PLE$ .*

We refer the reader to [Jeannerod et al. 2011] for a proof that any matrix has a PLE decomposition. There is also proven that the PLE decomposition is at least as efficient for almost all applications as other column rank profile revealing decompositions in terms of space and time complexity.

One last technicality stems from the fact that a computer’s memory is addressed in a one-dimensional fashion, while matrices are two-dimensional objects. The methods of this work use row major representations, not column major representations. The definition is as follows:

**DEFINITION 6 ROW MAJOR REPRESENTATION.** *A computer program stores a matrix in row major representation if the entire first row is stored in memory, followed by the entire second row, followed by the entire third row, and so forth.*

### 3. ALGORITHMS

The PLE decomposition can be computed in sub-cubic time complexity by a block recursive algorithm reducing most algebraic operations to matrix multiplication and some other much cheaper operations [Jeannerod et al. 2011]. We denote this algorithm as “*block recursive* PLE decomposition” in this work. In principle this recursion could continue until  $1 \times 1$  matrices are reached as a base case. However, in order to achieve optimal performance, these recursive calls must at some dimension “cross over” to a base case implementation, presumably much larger than  $1 \times 1$ . This base case has asymptotic complexity worse than the calling algorithm but better performance, in practice, for small dimensions.

In this work, we also focus on optimising this base case. This is done in two phases: a block iterative algorithm using tabulated updates (based on the “Method of four Russians” inversion algorithm [Bard 2006]) denoted by BLOCKITERATIVE-PLE on top of a lazy iterative partial decomposition which we denote PARTIALPLE. Here, lazy has a technical meaning which we define as follows. For some entry  $A[i, j]$




Fig. 1. Lazy PLE elimination: the matrix has been processed through row  $r$  and column  $c$ . In order to determine if the coefficient at position  $(i, j)$  is non zero, the algorithm must apply updates of the two last non zero columns as only  $s_0$  is greater than  $i$  and hence the update of its corresponding column was already applied.

below the main diagonal, i.e. with  $i > j$ , we do not modify row  $i$  in such a way as to make  $A[i, j] = 0$  until the last possible moment. As it comes to pass, that moment will be when we check to see if  $A[i, i] = 0$ , to see if it would be a suitable pivot element.

### 3.1 Partial iterative algorithm

This algorithm is used as a base case of the block iterative algorithm of section 3.2. It is based on the classical  $\mathcal{O}(n^3)$  Gaussian elimination algorithm, where pivots are found by searching column-wise: the algorithm only moves to the next column when all rows positions of the current column have been found to be zero. This condition ensures that the  $E$  matrix will be in row echelon form and the column indices of the leading entries of each row (i.e. the pivots) will indicate the column rank profile of the input matrix.

To improve cache efficiency, the updates are done in a “lazy” manner, inspired by the left looking elimination strategy: updates are delayed to the moment when the algorithm actually needs to check whether a coefficient is non-zero. At that moment, all trailing coefficients of the current row are also updated, as they are in the same cache page. This process requires the maintenance of a table  $s = [s_0, \dots, s_{r-1}]$  storing for each pivot column the highest-numbered row considered so far. It is defined by  $s_i = \max\{i_j, j = 0 \dots i\}$ , where  $i_j$  is the row number where the  $j$ th column’s pivot has been found. Consequently, the sequence  $s_i$  is monotonically decreasing. As a result, the algorithm produces the PLE decomposition of only  $s = \max\{i_j, j = 0 \dots i\}$  rows of the input matrix. We therefore call it a “lazy” iterative partial PLE decomposition. The complete algorithm is shown as Algorithm 1 and the key idea is illustrated in Figure 1.

As a specialized instance of standard Gaussian elimination, algorithm 1 computes a partial PLE decomposition of an  $m \times n$  matrix of rank  $r$  in  $\mathcal{O}(mnr)$  operations over  $\mathbb{F}_2$ . For the square case, this is obviously  $\mathcal{O}(n^2r)$ .

**Algorithm 1:** PARTIALPLE: lazy iterative partial PLE Decomposition**Input:**  $A$  – an  $m \times n$  matrix**Result:** Returns  $r$ , the rank of  $A$ , a permutation vector  $P$  of size  $s$ , a vector  $Q$  of size  $r$  and an integer  $s$  with  $r \leq s \leq m$ **Ensures:**  $A \leftarrow \begin{bmatrix} L \setminus E \\ A_1 \end{bmatrix}$  where  $L$  is  $s \times r$  unit lower triangular,  $E$  is an  $r \times n$  echelon form,  $A = \begin{bmatrix} A_0 \\ A_1 \end{bmatrix}$ ,  $PLE = A_0$  and  $\text{rank}(A_0) = r$ .

```

1 begin
2    $r, c \leftarrow 0, 0$ ;
   //  $n$  dim. array, s.t. column  $i$  has been updated to row  $s_i$ .
3    $s \leftarrow [0, \dots, 0]$ ;
4   while  $r < m$  and  $c < n$  do
5     found  $\leftarrow$  False;
6     for  $c \leq j < n$  do // search for some pivot
7       for  $r \leq i < m$  do
8         if  $r > 0$  and  $i > s_{r-1}$  then
9           Let  $k \leq r-1$  be such that  $s_{k-1} \geq i > s_k$ ;
10          for  $k \leq l < r$  do // Update row  $i$  from column  $j$  to  $n$ 
11            if  $A[i, Q_l] \neq 0$  then add the row  $l$  to the row  $i$  starting
              at column  $j$ ;
12             $s_l \leftarrow i$ ;
13          if  $A[i, j] \neq 0$  then found  $\leftarrow$  True and break;
14        if found then break;
15      if found then
16         $P_r, Q_r \leftarrow i, j$ ;
17        swap the rows  $r$  and  $i$ ;
18         $r, c \leftarrow r + 1, j + 1$ ;
19      else
20         $j \leftarrow j + 1$ ;
21     $s_{\max} \leftarrow \max(s_i)$ ;
22    for  $0 \leq j < r$  do
23      for  $s_j + 1 \leq i \leq s_{\max}$  do
24        if  $A[i, Q_j] \neq 0$  then
25          add the row  $j$  to the row  $i$  starting at column  $j$ ;
26    // Now compress  $L$ 
27    for  $0 \leq j < r$  do swap the columns  $j$  and  $Q_j$  starting at row  $j$ ;
    return  $P, r, Q, s$ ;

```

### 3.2 Block Iterative

For matrices over  $\mathbb{F}_2$  the “Method of the Four Russians” Inversion (M4RI), performs Gaussian elimination on a random  $n \times n$  matrix in  $\mathcal{O}(n^3/\log n)$ . The M4RI algorithm was introduced in [Bard 2006] and later described in [Bard 2007] and [Bard 2009, Ch. 9]. It inherits its name and main idea from “Konrod’s Method for matrix-matrix multiplication,” which is often mistakenly called “The Method of Four Russians” despite the fact that none of the authors of that paper are actually Russians [Arlazarov et al. 1970; Aho et al. 1974]. For the key idea behind the M4RI algorithm, assume that we have computed the row echelon form of an  $k \times n$  submatrix of some matrix  $A$  and that this submatrix has rank  $k$ . Then, in order to eliminate below, we may use a table of all  $2^k$  linear combinations of the  $k$  pivot rows found in the previous step. Hence, with only one addition per row,  $k$  columns are processed. Furthermore, the construction of all  $2^k$  linear combinations of the pivots can be accomplished in only  $2^k$  vector additions saving a factor of  $k$  compared to naive table construction. Finally, if we set  $k \approx \log n$  this algorithm can achieve  $\mathcal{O}(n^3/\log n)$  under the assumption that it always finds pivots among the first  $c \cdot k$  rows it considers where  $c$  is a small constant. If this condition is not met it either fails (see also [Bard 2007]), or regresses to cubic complexity (this is the strategy is implemented in the M4RI library). We refer the reader to Appendix A for a more in-depth treatment of the M4RI algorithm.

Our implementation of block-iterative PLE decomposition is based on our implementation of the M4RI algorithm and the performance reported in Section 5 is partly due to techniques inherited from our fast implementation of the M4RI algorithm. Thus, we point out that our implementation of the M4RI algorithm is very efficient, allowing it to beat asymptotically fast implementations of Gaussian elimination over  $\mathbb{F}_2$  up to matrix dimensions of several thousand despite having a worse asymptotic complexity. (See also Appendix A.)

Algorithm 2, which can be seen as the PLE analog of the M4RI algorithm<sup>3</sup>, does not require pivots among the first  $c \cdot k$  rows in order to achieve  $\mathcal{O}(n^3/\log n)$ .

Informally, the algorithm proceeds as follows. We divide the input matrix  $A$  into vertical stripes of  $k \approx \log n$  columns each. Then, we perform PLE decomposition on the first such stripe. This operation is cheap because these stripes are narrow. Now, in order to update the matrix to the right of the current stripe we use the M4RM matrix multiplication algorithm [Albrecht et al. 2010], i.e., we compute all possible linear combinations of the pivot rows just discovered and use the rows of the transformation matrix  $L$  as indices in this table. Hence, for each row of  $A_{22}$  we perform one row addition from the table instead of  $\mathcal{O}(k)$  row additions. Thus, regardless of the rank of the currently considered stripe, we achieve  $\mathcal{O}(n^3/\log n)$  because the update on the right is performed using M4RM.

Note, however, that our presentation of Algorithm 2 is independent of the choice of matrix-multiplication algorithm for clarity of presentation. Thus it could (in principle) be instantiated with naive multiplication, and that would result in a

<sup>3</sup>Note that [Albrecht 2010] and earlier versions of this work contain an algorithm denoted MMPF. As it turns out, the MMPF algorithm suffers from the same performance regression as M4RI, in the case when pivots cannot be found among the first  $c \cdot k$  considered rows. On the other hand, unlike M4RI, it does compute a PLE decomposition.

complexity of  $\mathcal{O}(n^3)$ .

However, performing PLE decomposition on the current slice has an important disadvantage when compared with the M4RI algorithm: it is less cache friendly. Recall, that our matrices are stored in row major representation. PLE decomposition will always touch all rows when treating the current slice. This is cheap when considering field operations only. However, if the matrix  $A$  is sufficient large, this will result in a load of a new cache line for each considered row. To mitigate this effect, we use PARTIALPLE as defined in the last section.

Every call to PARTIALPLE returns some  $\bar{r}$  and some  $s$ , which are the dimensions of the  $L$  matrix calculated by the PARTIALPLE call. Recall, the dimensions of  $A$  are  $m \times n$  and  $k$  is the parameter of our algorithm. Assume that a call to PARTIALPLE has returned with  $\bar{r} = k$  and some  $s < m$  and hence that we can decompose our matrix as follows

$$A = \begin{pmatrix} A_{00} & A_{01} & A_{02} \\ A_{10} & A_{11} = L_{11} \setminus E_{11} & A_{12} \\ A_{20} & A_{21} = L_{21} & A_{22} \\ A_{30} & A_{31} & A_{32} \end{pmatrix}$$

where  $A_{11}$  has  $r$  rows and  $A_{21}$  has  $s - r$  rows. We note that if  $\bar{r} < k$  we have  $s = m$  and hence  $A_{31}$  a matrix with zero rows, i.e., we only need to update  $A_{22}$  which reduces to matrix multiplication as discussed above. Hence, we may focus on the  $\bar{r} = k$  case in our discussion.

Algorithm 2 proceeds by first computing  $A_{12} \leftarrow (L_{11})^{-1} \cdot A_{12}$ . We then construct a table  $T$  of all  $2^k$  linear combinations of the rows of  $A_{12}$ . Using Gray codes or similar techniques these tables can be constructed using  $2^k$  vector additions only [Albrecht et al. 2010]. Using this table, we can implement the map  $v \rightarrow v \cdot A_{12}$  for all  $2^k$  possible vectors  $v$  as the table lookup  $T(v)$ . Furthermore, we can implement the map  $v \rightarrow v \cdot (E_{11}^{-1} \cdot A_{12})$  as  $(v \cdot E_{11}^{-1}) \cdot A_{12}$ . That is, we store a table  $T'$  of all  $2^k$  possible pairs  $(v, v \cdot E_{11}^{-1})$  and compute  $v \rightarrow v \cdot E_{11}^{-1} \cdot A_{12}$  by the table lookup  $T(T'(v))$ .

Finally, we “postprocess”  $T$  such that we also update  $L$  when computing  $v \cdot (E_{11})^{-1}$ , i.e., we replace any row  $v$  of  $A_{31}$  by  $v \cdot E_{11}^{-1}$  when adding  $(v \cdot E_{11}^{-1}) \cdot A_{12}$  to a row of  $A_{32}$ . For example, assume  $\bar{r} = 3$  and that the first row of the  $\bar{r} \times n$  submatrix  $A_{11}$  has the first  $\bar{r}$  bits equal to  $[1 \ 0 \ 1]$  ( $= 5$ ). Assume further that we want to clear  $\bar{r}$  bits of a row which also starts with  $[1 \ 0 \ 1]$ . Then – in order to generate  $L$  – we need to encode that this row is cleared by adding the first row only, i.e., we want the first  $\bar{r} = 3$  bits to be  $[1 \ 0 \ 0]$ . Thus, to correct the table, we add – in the postprocessing phase of  $T$  –  $[1 \ 0 \ 0]$  to the first  $\bar{r}$  bits  $T(5)$ . Hence, when adding this row of  $T$  to  $A_{31}$ , we produce  $[1 \ 0 \ 1] \oplus [1 \ 0 \ 1] \oplus [1 \ 0 \ 0] = [1 \ 0 \ 0]$  and store it in  $A_{31}$ .

Using  $T$  and  $T'$ , we can now perform the operations

$$A_{22} \leftarrow L_{21} \cdot A_{12} + A_{22} \text{ in line 14.}$$

$$(A_{31}, A_{32}) \leftarrow (A_{31} \cdot (E_{11})^{-1}, A_{31} \cdot (E_{11})^{-1} \cdot A_{12} + A_{32}) \text{ in line 15}$$

using one vector addition per row. The update of  $A_{22}$  is a multiplication by  $A_{12}$  plus an addition and is hence accomplished by one lookup to  $T$  and one vector



addition per row. Similarly, the updates of  $A_{31}$  and  $A_{32}$  can be achieved using lookups in  $T'$  and  $T$  and one vector addition per row.

---

**Algorithm 2:** Block-iterative PLE Decomposition

---

**Input:**  $A$  – an  $m \times n$  matrix  
**Result:** Returns  $r$ , the rank of  $A$ , a permutation  $P$  of size  $m$  and an  $r$ -dimensional vector  $Q$   
**Ensures:**  $A \leftarrow [L \setminus E]$  where  $L$  is unit lower triangular and  $E$  is in echelon form such that  $A = PLE$  and  $Q_i$  is the column index of the  $i$ th pivot of  $E$ .

```

1 begin
2    $r, c \leftarrow 0, 0$ ;
3    $k \leftarrow$  pick some  $0 < k \leq n$ ; //  $\mathcal{O}(n^3 / \log n) \rightarrow k \approx \log n$ 
4    $Q \leftarrow [0, \dots, 0]$ ;
5   while  $r < m$  and  $c < n$  do
6     if  $k > n - c$  then
7        $k \leftarrow n - c$ ;
8     /*  $A = \begin{bmatrix} A_{00} & A_{01} & A_{02} \\ A_{00} & A_{11} & A_{12} \end{bmatrix}$  where  $A_{00}$  is  $r \times c$  and  $A_{11}$  is  $(n - r) \times k$  */
9      $\bar{r}, \bar{P}, \bar{Q}, s \leftarrow \text{PARTIALPLE}(A_{11})$ ;
10    for  $0 \leq i < \bar{r}$  do
11       $Q_{r+i} \leftarrow c + \bar{Q}_i$ ;
12       $P_{r+i} \leftarrow r + \bar{P}_i$ ;
13      /*  $A = \begin{bmatrix} A_{00} & A_{01} & A_{02} \\ A_{10} & L_{11} \setminus E_{11} & A_{12} \\ A_{20} & L_{21} & 0 & A_{22} \\ A_{30} & A_{31} & A_{32} \end{bmatrix}$  where  $A_{00}$  is  $r \times c$  and  $L_{11}$  is
14          $\bar{r} \times \bar{r}$ ,  $E_{11}$  is  $\bar{r} \times k$  */
15       $\begin{bmatrix} A_{10} \\ A_{20} \end{bmatrix} \leftarrow P^T \begin{bmatrix} A_{10} \\ A_{20} \end{bmatrix}$ ,  $\begin{bmatrix} A_{12} \\ A_{22} \end{bmatrix} \leftarrow P^T \begin{bmatrix} A_{12} \\ A_{22} \end{bmatrix}$ ;
16       $A_{12} \leftarrow (L_{11})^{-1} \times A_{12}$ ;
17       $A_{22} \leftarrow L_{21} \cdot A_{12} + A_{22}$ ;
18       $(A_{31}, A_{32}) \leftarrow (A_{31} \cdot E_{11}^{-1}, A_{31} \cdot E_{11}^{-1} \cdot A_{12} + A_{32})$ ;
19      // Now compress L
20      for  $0 \leq i < \bar{r}$  do
21        swap the columns  $r + i$  and  $c + i$  starting at row  $r + i$ ;
22       $r, c \leftarrow r + \bar{r}, c + k$ ;
23 return  $P, Q, r$ ;
```

---

Overall we have:

LEMMA 1. *Let  $k = \log n$ , then Algorithm 2 computes the PLE decomposition of some  $n \times n$  matrix  $A$  in  $\mathcal{O}(n^3 / \log n)$  field operations if the M4RM algorithm is*

used in line 14 and 15.

PROOF. Let  $r_i$  be the rank computed by each call to Algorithm 1. Note that the sum of the  $r_i$ 's is the rank of the matrix,  $r$ .

The overall operation count is obtained the sum of the following contributions:

—calls to algorithm 1 (line 8):

$$\sum_{i=0}^{n/k-1} (n - \sum_{j=0}^{i-2} r_j) k r_i \leq n k r,$$

—solving with the  $r_i \times r_i$  lower triangular matrix  $L_{11}$  on  $A_{12}$  (line 13) can be done with the classic cubic time back substitution algorithm:

$$\sum_{i=0}^{n/k-1} r_i^2 (n - (i+1)k) \leq \sum_{i=0}^{n/k-1} k^2 n \leq n^2 k,$$

—updates based on Gray code tabulated matrix multiplication (line 14):

$$\sum_{i=0}^{n/k-1} (n - (i+1)k) (2^k + n - \sum_{j=0}^i r_j) \leq (2^k + n) \sum_{i=1}^{n/k} i k \leq \frac{(2^k + n) n^2}{k},$$

—updates in line 15 share the same table as the previous one, line 14 (all linear combinations of rows of  $A_{12}$ ). Hence they only take:

$$\sum_{i=0}^{n/k-1} (n - (i+1)k) r_i \leq \frac{r n^2}{k}.$$

For  $k = \log n$ , the overall complexity is  $\mathcal{O}(n^3 / \log n)$ .  $\square$

### 3.3 Block Recursive

In [Jeannerod et al. 2011] an algorithm is given that computes a PLE decomposition in-place and in time complexity  $\mathcal{O}(n^\omega)$  by reduction to matrix-matrix multiplication. This algorithm can be recovered by setting  $k = n/2$  in Algorithm 2 and by replacing PARTIALPLE with PLE. We give it explicitly in Algorithm 3 for clarity of presentation. The asymptotic complexity of the algorithm depends on matrix multiplication algorithm used. Our implementation uses Strassen-Winograd and hence achieves the  $\mathcal{O}(n^\omega)$  running time; this is the subject of a previous paper [Albrecht et al. 2010]. We note that  $A_{01} \leftarrow (L_{00})^{-1} \times A_{01}$  is the well-known triangular solving with matrices (TRSM) operation with a lower triangular matrix on the left which can be reduced to matrix-matrix multiplication as well [Jeannerod et al. 2011].

## 4. IMPLEMENTATION DETAILS

In our implementation we call Algorithm 3 which recursively calls itself until the matrix fits into L2 cache when it calls Algorithm 2. Our implementation of Algorithm 2 uses four Gray code tables. (See also [Albrecht et al. 2010].)

One of the major bottlenecks are column swaps. Note that this operation is usually not considered when estimating the complexity of algorithms in linear algebra,

**Algorithm 3:** Block-recursive PLE Decomposition**Input:**  $A$  – an  $m \times n$  matrix**Result:** Returns  $r$ , the rank of  $A$ , a permutation  $P$  of size  $m$  and an  $r$ -dimensional vector  $Q$ **Ensures:**  $A \leftarrow [L \setminus E]$  where  $L$  is unit lower triangular and  $E$  is in echelon form such that  $A = PLE$  and  $Q_i$  is the column index of the  $i$ th pivot of  $E$ .

```

1 begin
2    $n' \leftarrow$  pick some integer  $0 \leq n' < n$ ; //  $n' \approx n/2$ 
   /*  $A = \begin{bmatrix} A_0 & A_1 \end{bmatrix}$  where  $A_0$  is  $m \times n'$ . */
3    $r', P', Q' \leftarrow \text{PLE}(A_0)$ ; // first recursive call
4   for  $0 \leq i \leq r'$  do
5      $Q_i \leftarrow Q'_i$ ;
6   for  $0 \leq i \leq r'$  do
7      $P_i \leftarrow P'_i$ ;
   /*  $A = \begin{bmatrix} L_{00} \setminus E_{00} & A_{01} \\ A_{10} & A_{11} \end{bmatrix}$  where  $L_{00}$  is  $r' \times r'$  and  $E_{00}$  is  $r' \times n'$ . */
8   if  $r'$  then
9      $A_1 \leftarrow P \times A_1$ ;
10     $A_{01} \leftarrow (L_{00})^{-1} \times A_{01}$ ;
11     $A_{11} \leftarrow A_{11} + A_{10} \times A_{01}$ ;
12    $r'', P'', Q'' \leftarrow \text{PLE}(A_{11})$ ; // second recursive call
13    $A_{10} \leftarrow P'' \times A_{10}$ ;
   /* Update P & Q */
14   for  $0 \leq i < m - r'$  do
15      $P_{r'+i} \leftarrow P''_i + r'$ ;
16   for  $0 \leq i < n - n'$  do
17      $Q_{n'+i} \leftarrow Q''_i + n'$ ;
18    $j \leftarrow r'$ ;
19   for  $n' \leq i < n' + r''$  do
20      $Q_j \leftarrow Q_i$ ;  $j \leftarrow j + 1$ ;
   /* Now compress L */
21    $j \leftarrow n'$ ;
22   for  $r' \leq i < r' + r''$  do
23     swap the columns  $i$  and  $j$  starting at row  $i$ ;
24   return  $P, Q, r' + r''$ ;
```

since it is not a field operation. In Algorithm 4 a simple algorithm for swapping two columns  $a$  and  $b$  is given with bit-level detail. In Algorithm 4 we assume that the bit position of  $a$  is greater than the bit position of  $b$  for clarity of presentation. The advantage of the strategy in Algorithm 4 is that it uses no conditional jumps in the inner loop. However, it still requires 9 instructions per row. On the other hand, we can add two rows with  $9 \cdot 128 = 1152$  entries in 9 instructions if the “SSE2 instruction set” is available. Thus, for matrices of size  $1152 \times 1152$  it takes roughly the same number of instructions to add two matrices as it does to swap two columns. If we were to swap every column with one other column once during some algorithm it thus would be as expensive as a matrix multiplication (for matrices of these dimensions).

---

**Algorithm 4:** Column Swap

---

**Input:**  $A$  – an  $m \times n$  matrix  
**Input:**  $a$  – an integer  $0 \leq a < b < n$   
**Input:**  $b$  – an integer  $0 \leq a < b < n$   
**Result:** Swaps the columns  $a$  and  $b$  in  $A$

```

1 begin
2    $M \leftarrow$  the memory where  $A$  is stored;
3    $a_w, b_w \leftarrow$  the word index of  $a$  and  $b$  in  $M$ ;
4    $a_b, b_b \leftarrow$  the bit index of  $a$  and  $b$  in  $a_w$  and  $b_w$ ;
5    $\Delta \leftarrow a_b - b_b$ ;
6    $a_m \leftarrow$  the bit-mask where only the  $a_b$ -th bit is set to 1;
7    $b_m \leftarrow$  the bit-mask where only the  $b_b$ -th bit is set to 1;
8   for  $0 \leq i < m$  do
9      $R \leftarrow$  the memory where the row  $i$  is stored;
10     $R[a_w] \leftarrow R[a_w] \oplus ((R[b_w] \odot b_m) \gg \Delta)$ ;
11     $R[b_w] \leftarrow R[b_w] \oplus ((R[a_w] \odot a_m) \ll \Delta)$ ;
12     $R[a_w] \leftarrow R[a_w] \oplus ((R[b_w] \odot b_m) \gg \Delta)$ ;

```

---

Another bottleneck for relatively sparse matrices in dense “row-major representation” is the search for pivots. Searching for a non-zero element in a row can be relatively expensive due to the need to identify the bit position; hence, we cannot work with machine words in bulk but have to read individual bits and see which column they are in. However, the main performance penalty is due to the fact that searching for a non-zero entry in one column in a row-major representation is very cache unfriendly.

Indeed, both our implementation and the implementation available in MAGMA suffer from performance regression on relatively sparse matrices as shown in Figure 2. We stress that this is despite the fact that the theoretical complexity of matrix decomposition is rank sensitive. More precisely, strictly fewer field operations must be performed for low-rank matrices. This highlights the distinction between minimizing the running time and minimizing the number of field operations—these are not identical objectives. While the penalty for relatively sparse matrices is

much smaller for our implementation than for MAGMA, it clearly does not achieve the theoretically optimum performance.




Fig. 2. Gaussian elimination of  $10,000 \times 10,000$  matrices on Intel 2.33GHz Xeon E5345 comparing Magma 2.17-12 and M4RI 20111004.

## 5. PERFORMANCE

In Table I, we consider reduced row echelon forms, i.e., we first compute the PLE decomposition which we then use to recover the reduced row echelon form, of dense random matrices over  $\mathbb{F}_2$ . We compare the running times for Algorithm 3 (with either cubic PLE decomposition or Algorithm 2 as base case).

Matrix Dimension	base case: cubic PLE	base case: Algorithm 2
$10,000 \times 10,000$	1.478s	0.880s
$16,384 \times 16,384$	5.756s	3.570s
$20,000 \times 20,000$	8.712s	5.678s
$32,000 \times 32,000$	29.705s	24.240s

Table I. Comparing base cases on 64-bit Debian/GNU Linux, 2.33Ghz Core 2 Duo.

In Table II, we give the average running time over ten trials for computing reduced row echelon forms of dense random  $n \times n$  matrices over  $\mathbb{F}_2$ . We compare the asymptotically fast implementation due to Allan Steel in MAGMA, the cubic Gaussian elimination implemented by Victor Shoup in NTL, and both our implementations. Both the implementation in MAGMA and our PLE decomposition reduce matrix decomposition to matrix multiplication. A discussion and comparison of matrix multiplication in the M4RI library and in MAGMA can be found in

[Albrecht et al. 2010]. In Table II, the column ‘PLE’ denotes the complete running time for first computing the PLE decomposition and the computation of the reduced row echelon form from the PLE form.

64-bit Linux, 2.6Ghz Opteron				
$n$	MAGMA 2.15-10	NTL 5.4.2	M4RI 20090105	PLE 20100324
10,000	3.351s	18.45s	2.430s	1.452s
16,384	11.289s	72.89s	10.822s	6.920s
20,000	16.734s	130.46s	19.978s	10.809s
32,000	57.567s	479.07s	83.575s	49.487s
64,000	373.906s	2747.41s	537.900s	273.120s
64-bit Linux, 2.33Ghz Xeon (E5345)				
$n$	MAGMA 2.16-7	NTL 5.4.2	M4RI 20100324	PLE 20100324
10,000	2.660s	12.05s	1.360s	0.864s
16,384	8.617s	54.79s	5.734s	3.388s
20,000	12.527s	100.01s	10.610s	5.661s
32,000	41.770s	382.52s	43.042s	20.967s
64,000	250.193s	—	382.263s	151.314s

Table II. Gaussian elimination for random matrices  $n \times n$  matrices

## REFERENCES

- AHO, A., HOPCROFT, J., AND ULLMAN, J. 1974. *The Design and Analysis of Computer Algorithms*. Addison-Wesley.
- ALBRECHT, M., BARD, G., AND HART, W. 2010. Algorithm 898: Efficient multiplication of dense matrices over  $\text{GF}(2)$ . *ACM Trans. Math. Softw.* 37, 9:1–9:14.
- ALBRECHT, M. R. 2010. Algorithmic Algebraic Techniques and their application to Block Cipher Cryptanalysis. Ph.D. thesis, Royal Holloway, University of London.
- ALBRECHT, M. R. AND BARD, G. V. 2011. *The M4RI Library – Version 20111004*. The M4RI Team. <http://m4ri.sagemath.org>.
- ARLAZAROV, V., DINIC, E., KRONROD, M., AND FARADZEV, I. 1970. On economical construction of the transitive closure of a directed graph. *Dokl. Akad. Nauk.* 194, 11. (in Russian), English Translation in Soviet Math Dokl.
- BARD, G. V. 2006. Accelerating cryptanalysis with the Method of Four Russians. Cryptology ePrint Archive, Report 2006/251. Available at <http://eprint.iacr.org/2006/251.pdf>.
- BARD, G. V. 2007. Algorithms for solving linear and polynomial systems of equations over finite fields with applications to cryptanalysis. Ph.D. thesis, University of Maryland.
- BARD, G. V. 2009. *Algebraic Cryptanalysis*. Springer Verlag, Berlin, Heidelberg, New York.
- BOSMA, W., CANNON, J., AND PLAYOUST, C. 1997. The MAGMA Algebra System I: The User Language. In *Journal of Symbolic Computation* 24. Academic Press, 235–265.
- COPPERSMITH, D. AND WINOGRAD, S. 1990. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation* 9, 3, 251 – 280. `journal|title|Computational algebraic complexity editorial|ce:title|`.
- COURTOIS, N. T., KLIMOV, A., PATARIN, J., AND SHAMIR, A. 2000. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In *Advances in Cryptology — EUROCRYPT 2000*. Lecture Notes in Computer Science, vol. 1807. Springer Verlag, Berlin, Heidelberg, New York, 392–407.
- FAUGÈRE, J.-C. 1999. A new efficient algorithm for computing Gröbner basis (F4). *Journal of Pure and Applied Algebra* 139, 1-3, 61–88.
- GOLUB, G. AND VAN LOAN, C. 1996. *Matrix Computations*, third ed. The Johns Hopkins University Press.
- IBARRA, O., MORAN, S., AND HUI, R. 1982. A generalization of the fast LUP matrix decomposition algorithm and applications. *Journal of Algorithms* 3, 45–56.
- JEANNEROD, C.-P., PERNET, C., AND STORJOHANN, A. 2011. Rank profile revealing Gaussain elimination and the CUP matrix decomposition. *pre-print to appear*.
- LAZARD, D. 1983. Gröbner-bases, Gaussian elimination and resolution of systems of algebraic equations. In *Proceedings of the European Computer Algebra Conference on Computer Algebra*. Lecture Notes in Computer Science, vol. 162. Springer Verlag, Berlin, Heidelberg, New York.

## A. THE M4RI ALGORITHM

Consider the matrix  $A$  of dimension  $m \times n$  in Figure 3. Because  $k = 3$  here, the  $3 \times n$  submatrix on the top has full rank and we performed Gaussian elimination on it. Now, we need to clear the first  $k$  columns of  $A$  in the rows below  $k$  (and above the submatrix in general if we want the reduced row echelon form). There are  $2^k$  possible linear combinations of the first  $k$  rows, which we store in a table  $T$ . We index  $T$  by the first  $k$  bits (e.g.  $011 \rightarrow 3$ ). Now to clear  $k$  columns of row  $i$  we use the first  $k$  bits of that row as an index in  $T$  and add the matching row of  $T$  to row  $i$ , causing a cancellation of  $k$  entries. Instead of up to  $k$  additions, this only costs one addition, due to the pre-computation. Using Gray codes (see also [Albrecht et al. 2010]) or similar techniques this pre-computation can be performed in  $2^k$  vector additions and the overall cost is  $2^k + m - k + k^2$  vector additions in the worst case (where  $k^2$  accounts for the Gauss elimination of the  $k \times n$  submatrix). The naive

approach would cost  $k \cdot m$  row additions in the worst case to clear  $k$  columns. If we set  $k = \log m$  then the complexity of clearing  $k$  columns is  $\mathcal{O}(m + \log^2 m)$  vector additions in contrast to  $\mathcal{O}(m \cdot \log m)$  vector additions using the naive approach.

$$A = \left( \begin{array}{ccc|cccc} 1 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & \dots \\ 0 & 1 & 0 & 1 & 1 & 1 & 0 & \dots \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 1 & \dots \\ \dots & & & & & & & & \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & \dots \\ \mathbf{1} & \mathbf{1} & \mathbf{0} & 0 & 1 & 0 & 1 & 1 & \dots \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & \dots \\ \dots & & & & & & & & \\ \mathbf{1} & \mathbf{1} & \mathbf{0} & 1 & 1 & 1 & 0 & 1 & \dots \end{array} \right) T = \left[ \begin{array}{ccc|cccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & \dots \\ 0 & 1 & 0 & 1 & 1 & 1 & 0 & \dots \\ 0 & 1 & 1 & 1 & 1 & 0 & 0 & \dots \\ 1 & 0 & 0 & 1 & 0 & 1 & 1 & \dots \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 & \dots \\ \mathbf{1} & \mathbf{1} & \mathbf{0} & 0 & 1 & 0 & 0 & \dots \\ 1 & 1 & 1 & 0 & 1 & 1 & 1 & \dots \end{array} \right]$$

Fig. 3. M4RI Idea

This idea leads to Algorithm 8. In this algorithm the subroutine GAUSSSUBMATRIX (see Algorithm 7) performs Gauss elimination on a  $k \times n$  submatrix of  $A$  starting at position  $(r, c)$  and searches for pivot rows up to  $m$ . If it cannot find a submatrix of rank  $k$  it will terminate and return the rank  $\bar{k}$  found so far.

The subroutine MAKETABLE (see Algorithm 6) constructs the table  $T$  of all  $2^k$  linear combinations of the  $k$  rows starting at row  $r$  and column  $c$ . More precisely, it enumerates all elements of the vector space  $\text{span}(r, \dots, r + \bar{k} - 1)$  spanned by the rows  $r, \dots, r + \bar{k} - 1$ . Finally, the subroutine ADDROWSFROMTABLE (see Algorithm 5) adds the appropriate row from  $T$  — indexed by  $k$  bits starting at column  $c$  — to each row of  $A$  with index  $i \notin \{r, \dots, r + \bar{k} - 1\}$ . That is, it adds the appropriate linear combination of the rows  $\{r, \dots, r + \bar{k} - 1\}$  onto a row  $i$  in order to clear  $k$  columns.

---

**Algorithm 5:** ADDROWSFROMTABLE

---

**Input:**  $A$  — an  $m \times n$  matrix  
**Input:**  $r_{\text{start}}$  — an integer  $0 \leq r_{\text{start}} < m$   
**Input:**  $r_{\text{end}}$  — an integer  $0 \leq r_{\text{start}} \leq r_{\text{end}} < m$   
**Input:**  $c_{\text{start}}$  — an integer  $0 \leq c_{\text{start}} < n$   
**Input:**  $k$  — an integer  $k > 0$   
**Input:**  $T$  — a  $2^k \times n$  matrix  
**Input:**  $L$  — an integer array of length  $2^k$

```

1 begin
2   for  $r_{\text{start}} \leq i < r_{\text{end}}$  do
3      $id = \sum_{j=0}^k A_{i, c_{\text{start}}+j} \cdot 2^{k-j-1};$ 
4      $j \leftarrow L_{id};$ 
5     add row  $j$  from  $T$  to the row  $i$  of  $A$  starting at column  $c_{\text{start}};$ 

```

---



**Algorithm 6: MAKETABLE**


---

**Input:**  $A$  – an  $m \times n$  matrix  
**Input:**  $r_{\text{start}}$  – an integer  $0 \leq r_{\text{start}} < m$   
**Input:**  $c_{\text{start}}$  – an integer  $0 \leq c_{\text{start}} < n$   
**Input:**  $k$  – an integer  $k > 0$   
**Result:** Returns an  $2^k \times n$  matrix  $T$

```

1 begin
2    $T \leftarrow$  the  $2^k \times n$  zero matrix;
3   for  $1 \leq i < 2^k$  do
4      $j \leftarrow$  the row index of  $A$  to add according to the Gray code;
5     add row  $j$  of  $A$  to the row  $i$  of  $T$  starting at column  $c_{\text{start}}$ ;
6    $L \leftarrow$  integer array allowing to index  $T$  by  $k$  bits starting at column  $c_{\text{start}}$ ;
7   return  $T, L$ ;

```

---

**Algorithm 7: GAUSSSUBMATRIX**


---

**Input:**  $A$  – an  $m \times n$  matrix  
**Input:**  $r$  – an integer  $0 \leq r < m$   
**Input:**  $c$  – an integer  $0 \leq c < n$   
**Input:**  $k$  – an integer  $k > 0$   
**Input:**  $r_{\text{end}}$  – an integer  $0 \leq r \leq r_{\text{end}} < m$   
**Result:** Returns the rank  $\bar{k} \leq k$  and puts the  $\bar{k} \times (n - c)$  submatrix starting at  $A[r, c]$  in reduced row echelon form.

```

1 begin
2    $r_{\text{start}} \leftarrow r$ ;
3   for  $c \leq j < c + k$  do
4      $\text{found} \leftarrow \text{False}$ ;
5     for  $r_{\text{start}} \leq i < r_{\text{end}}$  do
6       for  $0 \leq l < j - c$  do // clear the first columns
7         if  $A[i, c + l] \neq 0$  then add row  $r + l$  to row  $i$  of  $A$  starting at
          column  $c + l$ ;
8       if  $A[i, j] \neq 0$  then // pivot?
9         Swap the rows  $i$  and  $r_{\text{start}}$  in  $A$ ;
10        for  $r \leq l < r_{\text{start}}$  do // clear above
11          if  $A[l, j] \neq 0$  then add row  $r_{\text{start}}$  to row  $l$  in  $A$  starting at
            column  $j$ ;
12         $r_{\text{start}} \leftarrow r_{\text{start}} + 1$ ;
13         $\text{found} \leftarrow \text{True}$ ;
14        break;
15     if  $\text{found} = \text{False}$  then
16       return  $j - c$ ;
17   return  $j - c$ ;

```

---

**Algorithm 8: M4RI**


---

**Input:**  $A$  – an  $m \times n$  matrix  
**Input:**  $k$  – an integer  $k > 0$   
**Result:**  $A$  is in reduced row echelon form.

```

1 begin
2    $r, c \leftarrow 0, 0;$ 
3   while  $c < n$  do
4     if  $c + k > n$  then  $k \leftarrow n - c;$ 
5      $\bar{k} \leftarrow \text{GAUSSSUBMATRIX}(A, r, c, k, m);$ 
6     if  $\bar{k} > 0$  then
7        $T, L \leftarrow \text{MAKETABLE}(A, r, c, \bar{k});$ 
8        $\text{ADDEROWSFROMTABLE}(A, 0, r, c, \bar{k}, T, L);$ 
9        $\text{ADDEROWSFROMTABLE}(A, r + \bar{k}, m, c, \bar{k}, T, L);$ 
10     $r, c \leftarrow r + \bar{k}, c + \bar{k};$ 
11    if  $k \neq \bar{k}$  then  $c \leftarrow c + 1;$ 
  
```

---

When studying the performance of Algorithm 8, we expect the function MAKE-TABLE to contribute most. Instead of performing  $\bar{k}/2 \cdot 2^{\bar{k}} - 1$  additions MAKE-TABLE only performs  $2^{\bar{k}} - 1$  vector additions. However, in practice the fact that  $\bar{k}$  columns are processed in each loop iteration of ADDROWSFROMTABLE contributes significantly due to the better cache locality. Assume the input matrix  $A$  does not fit into L2 cache. Gaussian elimination would load a row from memory, clear one column and likely evict that row from cache in order to make room for the next few rows before considering it again for the next column. In the M4RI algorithm more columns are cleared per memory load instruction.

We note that our presentation of M4RI differs somewhat from that in [Bard 2007]. The key difference is that our variant does not throw an error if it cannot find a pivot within the first  $3k$  rows in GAUSSSUBMATRIX. Instead, our variant searches all rows and consequently the worst-case complexity is cubic. However, on average for random matrices we expect to find a pivot within  $3k$  rows with overwhelming probability [Bard 2009, Ch. 9.5.3] and thus the average-case complexity is  $\mathcal{O}(n^3/\log n)$ .

In the Table III, we give running times for computing the reduced row echelon form (or row echelon form for NTL) for random matrices of various dimensions. The column ‘M4RI’ in both tables refers to our implementation of Algorithm 8 using implementation techniques from [Albrecht et al. 2010] such as multiple pre-computation tables. The implementation that NTL uses is straight-forward Gaussian elimination and thus serves as a baseline for comparison while MAGMA implements asymptotically fast matrix elimination.

Table III shows that our implementation of the M4RI algorithm (with average-case complexity  $\mathcal{O}(n^3/\log n)$ ) is competitive with MAGMA’s asymptotically fast algorithm (with complexity  $\mathcal{O}(n^\omega)$ ) up to  $64,000 \times 64,000$  on the Core 2 Duo and up to  $16,384 \times 16,384$  on the Opteron.

64-bit Debian/GNU Linux, 2.33Ghz Core 2 Duo			
Matrix Dimensions	MAGMA 2.14-17	NTL 5.4.2	M4RI 20090105
$10,000 \times 10,000$	2.474s	14.140s	1.532s
$16,384 \times 16,384$	7.597s	67.520s	6.597s
$20,000 \times 20,000$	13.151s	123.700s	12.031s
$32,000 \times 32,000$	39.653s	462.320s	40.768s
$64,000 \times 64,000$	251.346s	2511.540s	241.017s
64-bit Debian/GNU Linux, 2.6Ghz Opteron			
Matrix Dimensions	MAGMA 2.14-13	NTL 5.4.2	M4RI 20090105
$10,000 \times 10,000$	3.351s	18.45s	2.430s
$16,384 \times 16,384$	11.289s	72.89s	10.822s
$20,000 \times 20,000$	16.734s	130.46s	19.978s
$32,000 \times 32,000$	57.567s	479.07s	83.575s
$64,000 \times 64,000$	373.906s	2747.41s	537.900s

Table III. RREF for Random Matrices